Spectral Partitioning Works: Planar graphs and finite element meshes

Preliminary Draft

Daniel A. Spielman*　　Shang-Hua Teng†

February 13, 1996

Abstract

Spectral partitioning methods use the Fiedler vector—the eigenvector of the second-smallest eigenvalue of the Laplacian matrix—to find a small separator of a graph. These methods are important components of many scientific numerical algorithms and have been demonstrated by experiment to work extremely well. In this paper, we show that spectral partitioning methods work well on bounded-degree planar graphs and finite element meshes—the classes of graphs to which they are usually applied. While naive spectral bisection does not necessarily work, we prove that spectral partitioning techniques can be used to produce separators whose ratio of vertices removed to edges cut is $O(\sqrt{n})$ for bounded-degree planar graphs and two-dimensional meshes and $O\left(n^{1/d}\right)$ for well-shaped d-dimensional meshes. The heart of our analysis is an upper bound on the second-smallest eigenvalues of the Laplacian matrices of these graphs.

1. Introduction

Spectral partitioning has become one of the most successful heuristics for partitioning graphs and matrices. It is used in many scientific numerical applications, such as mapping finite element calculations on parallel machines [Sim91, Wil90], solving sparse linear systems [PSW92], and partitioning for domain decomposition [CR87, CS93]. It is also used in VLSI circuit design and simulation [CSZ93, HK92, AK95]. Substantial experimental work has demonstrated that spectral methods find good partitions of the graphs and matrices that arise in many applications [BS92, HL92, HL93, PSL90, Sim91, Wil90]. However, the quality of the partition that

*Computer Science Division, U.C. Berkeley, CA 94720, spielman@cs.berkeley.edu. Supported in part by an NSF postdoc. After July 1996, Department of Mathematics, M.I.T.

†Department of Computer Science, University of Minnesota, Minneapolis, MN 55455, steng@cs.umn.edu. Supported in part by an NSF CAREER award (CCR-9502540) and an Intel research grant. Part of the work was done while visiting Xerox PARC.

these methods should produce has so far eluded precise analysis. In this paper, we will prove that spectral partitioning methods give good separators for the graphs to which they are usually applied.

The size of the separator produced by spectral methods can be related to the Fiedler value—the second smallest eigenvalue of the Laplacian—of the adjacency structure to which they are applied. By showing that well-shaped meshes in d dimensions have Fiedler value at most $O\left(1/n^{2/d}\right)$, we show that spectral methods can be used to find bisectors of these graphs of size at most $O\left(n^{1-1/d}\right)$. While a small Fiedler value does not immediately imply that there is a cut along the Fiedler vector that is a balanced separator, it does mean that there is a cut whose ratio of vertices separated to edges cut is $O\left(n^{1/d}\right)$. By removing the vertices separated by this cut, computing a Fiedler vector of the new graph, and iterating as necessary, one can find a bisector of $O\left(n^{1-1/d}\right)$ edges. In particular, we prove that bounded-degree planar graphs have Fiedler value at most $O(1/n)$, which implies that spectral techniques can be used to find bisectors of size at most $O(\sqrt{n})$ in these graphs. These bounds are the best possible for well-shaped meshes and planar graphs.

1.1. History

The spectral method of graph partitioning was born in the works of Donath and Hoffman [DH72, DH73] who first suggested using the eigenvectors of adjacency matrices of graphs to find partitions. Fiedler [Fie73, Fie75a, Fie75b] associated the second-smallest eigenvalue of the Laplacian of a graph with its connectivity and suggested partitioning by splitting vertices according to their value in the corresponding eigenvector. Thus, we call this eigenvalue the *Fiedler value* and a corresponding vector a *Fiedler vector*.

A few years later, Barnes and Hoffman [Bar82, BH84] used linear programming in combination with an examination of the eigenvectors of the adjacency matrix of a graph. In a similar vein, Boppana [Bop87] analyzed eigenvector techniques in conjunction with convex programming. However, the use of linear and convex programming made these techniques impractical for most applications.

By recognizing a relation between the Fiedler value and the Cheeger constant [Che70] of continuous manifolds, Alon [Alo86] and Sinclair and Jerrum [SJ89] demonstrated that if the Fiedler value of a graph is small, then directly partitioning the graph according to the values of vertices in the eigenvector will produce a cut with a good ratio of cut edges to separated vertices (see also [AM85, Fil91, DS91, Mih89, Moh89]). Around the same time, improvements in algorithms for approximately computing eigenvectors, such as the Lanczos algorithm, made the computation of eigenvectors practical [PSS82, Sim91]. In the next few years, a wealth of experimental work demonstrated that spectral partitioning methods work well on graphs that usually arise in practice [BS92, HL92, PSL90, Sim91, Wil90]. Spectral partitioning became a standard tool for mesh partitioning in many areas [HL93]. Still, researchers were unable to prove that spectral partitioning techniques would work well on the graphs encountered in practice. This failure is partially explained by results of Guattery and Miller [GM95] demonstrating that naive applications of spectral partitioning, such as spectral bisection, will fail miserably on some graphs that could conceivably arise in practice. By bounding the Fiedler values of the graphs of interest in scientific applications—bounded-degree planar graphs and well-shaped meshes—we are able to show that spectral partitioning methods will successfully find good partitions of these graphs.

In a related line of research, algorithms were developed along with proofs that they will always find small separators in various families of graphs. The seminal work in this area was that of Lipton and Tarjan [LT79], who constructed a linear-time algorithm that produces a 1/3-separator of $\sqrt{8n}$ nodes in any n-node planar graph. Their result improved a theorem of Ungar [Ung51] which demonstrated that every planar graph has a separator of size $O(\sqrt{n}\log n)$. Gilbert, Hutchinson, and Tarjan [GHT84] extended these results to show that every graph of genus at most g has a separator of size $O\left(\sqrt{gn}\right)$. Another generalization was obtained by Alon, Seymour, and Thomas [AST90], who showed that graphs that do not have an h-clique minor have separators of $O(h^{3/2}\sqrt{n})$ nodes. Plotkin, Rao, and Smith [PRS94] reduced the dependency on h from $h^{3/2}$ to h. Using geometric techniques, Miller, Teng, Thurston, and Vavasis [MT90, MTTV96a, MTTV96b, MTV91, MV91, Ten91] extended the planar separator theorem to graphs embedded in higher dimensions and showed that every well-shaped mesh in R^d has a $1/(d+2)$-separator of size $O(n^{1-1/d})$. Using multicommodity flow, Leighton and Rao [LR88] designed a partitioning method guaranteed to return a cut whose ratio of cut size to vertices separated is within logarithmic factors of optimal. While spectral methods have been favored in practice, they lacked a proof of effectiveness.

1.2. Outline of paper

In Section 2, we introduce the concept of a graph partition, review some facts from linear algebra that we require, and describe the class of spectral partitioning methods.

In Section 3, we prove the *embedding lemma*, which relates the quality of geometric embeddings of a graph with its Fiedler value. We then show (using the main result of Section 4) that every planar graph has a "nice" embedding as a collection of spherical caps on the surface of a unit sphere in three dimensions. By applying the embedding lemma to this embedding, we prove that the Fiedler value of every bounded-degree planar graph is $O(1/n)$.

In Section 4, we show that, for almost every arrangement of spherical caps on the unit sphere in R^d, there is a sphere-preserving map that transforms the caps so that the center of the sphere is the centroid of their centers. It is this fact that enables us to find nice embeddings of planar graphs.

In Sections 5 and 6, we extend our spectral planar separator theorem to the class of overlap graphs of k-ply neighborhood systems embedded in any fixed dimension. This extension enables us to show that the spectral method finds cuts of ratio $O(1/n^{1/d})$ for k-nearest neighbor graphs and well-shaped finite element meshes.

In Section 7, we present an elementary proof that from any vector perpendicular to the all-ones vector with small Rayleigh quotient, one can obtain a cut of small ratio.

In Section 8, we extend the results of Guattery and Miller to show that our results are essentially the best possible given current characterizations of well-shaped meshes. We present natural families of graphs for which Fiedler vectors can be used to find cuts of good ratio, but not good balance. We discuss why these graphs exist and why they might not appear in practice.

2. Introduction to Spectral Partitioning

In this section, we define the spectral partitioning method and introduce the terminology that we will use throughout the paper.

2.1. Graph Partitioning

Throughout this paper, $G = (V, E)$ will be a connected, undirected graph on n vertices.

A *partition* of a graph G is a division of its vertices into two disjoint subsets, A and $\bar{A}$. Without loss of generality, we can assume that $|A| \leq \left|\bar{A}\right|$. Let $E(A, \bar{A})$ be the set of edges with one endpoint in A and the other in $\bar{A}$. The *cut size* of the partition $(A, \bar{A})$ is simply $|E(A, \bar{A})|$. The *cut ratio*, or simply the *ratio* of the cut, denoted $\phi(A, \bar{A})$, is equal to the ratio of the size of the cut to the size of A, namely,

$$\phi(A, \bar{A}) = \frac{|E(A, \bar{A})|}{\min(|A|, |\bar{A}|)}.$$

The isoperimetric number of a graph, which measures how good a ratio cut one can hope to find, is defined to be

$$\phi(G) = \min_{|A| \leq n/2} \frac{\left|E(A, \bar{A})\right|}{|A|}.$$

In Section 7, we describe a relation between the isoperimetric number of a graph and its Fiedler value.

A partition is a *bisection* of G if A and $\bar{A}$ differ in size by at most 1. For δ in the range $0 < \delta \leq 1/2$, a partition is called a δ*-separator* if $\min(|A|, |\bar{A}|) \geq \delta n$. We use the word *cut* to refer to a partition separating any number of vertices and reserve the word *separator* for partitions that are δ-separators for some $\delta > 0$. Given an algorithm that can find cuts of ratio ϕ in G and its subgraphs, we can find a bisector of G of size $O(\phi n)$ (see Lemma 22).

2.2. Laplacians and Fiedler Vectors

The adjacency matrix, $A(G)$, of a graph G is the $n \times n$ matrix whose (i, j)-th entry is 1 if $(i, j) \in E$ and 0 otherwise. The diagonal entries are defined to be 0. Let D be the $n \times n$ diagonal matrix with entries $D_{i,i} = d_i$, where d_i is the degree of the ith vertex of G. The *Laplacian*, $L(G)$, of the graph G is defined to be $L(G) = D - A$.

Let M be an $n \times n$ matrix. An n-dimensional vector $\vec{x}$ is an *eigenvector* of M if there is a scalar λ such that $M\vec{x} = \lambda\vec{x}$. λ is the *eigenvalue* of M corresponding to the eigenvector $\vec{x}$. If M is a real symmetric matrix, then all of its n eigenvalues are real. The only matrices we consider in this paper will be the Laplacians of graphs. Notice that the all-ones vector is an eigenvector of any Laplacian matrix and that its associated eigenvalue is 0. Because Laplacian matrices are positive semidefinite, all the other eigenvalues must be non-negative. We will focus on the second smallest eigenvalue, λ_2, of the Laplacian and an associated eigenvector $\vec{u}$. Fiedler called this eigenvalue the "algebraic connectivity of a graph", so we will call it the *Fiedler value* and an associated eigenvector a *Fiedler vector*.

The following properties of Fiedler values and vectors play an important role in this paper:

- The Fiedler value of a graph is greater than zero if and only if the graph is connected.
- A Fiedler vector $\vec{u} = (u_1, ..., u_n)$ satisfies

$$\sum_{i=1}^{n} u_i = 0,$$

  because all-ones vector is an eigenvector of the Laplacian and the eigenvectors of a symmetric matrix are orthogonal.

- The Fiedler value, λ_2, of G satisfies

$$\lambda_2 = \min_{\vec{x} \perp (1,1,...,1)} \frac{\vec{x}^T L(G) \vec{x}}{\vec{x}^T \vec{x}},$$

  with the minimum occurring only when $\vec{x}$ is a Fiedler vector.

- For any vector $\vec{x} \in R^n$, we have

$$\vec{x}^T L(G) \vec{x} = \sum_{(i,j) \in E} (x_i - x_j)^2.$$

Let M be a symmetric $n \times n$ matrix and $\vec{x}$ be an n-dimensional vector. Then, the *Rayleigh quotient* of $\vec{x}$ with respect to M is

$$\frac{\vec{x}^T M \vec{x}}{\vec{x}^T \vec{x}}.$$

For proofs of these statements and many other fascinating facts about the eigenvalues and eigenvectors of graphs consult one of [Str88, Moh88, CDS90, Big93].

The Fiedler value, λ_2, of a graph is closely linked to its isoperimetric number. If G is a graph on more than three nodes, then one can show [AM85, Moh89, SJ88]

$$\frac{\lambda_2}{2} \leq \phi(G) \leq \sqrt{\lambda_2(2d - \lambda_2)}.$$

In Section 7, we will focus on the second inequality.

2.3. Spectral Partitioning Methods

Let $\vec{u} = (u_1, ..., u_n)$ be a Fiedler vector of the Laplacian of a graph G. The idea of spectral partitioning is to find a *splitting value* s and partition the vertices of G into the set of i such that $u_i > s$ and the set such that $u_i \leq s$. We call such a partition a *Fiedler cut*. There are several popular choices for the splitting value s:

- **bisection**: s is the median of $\{u_1, ..., u_n\}$.

- **ratio cut**: s is the value that gives the best ratio cut.
- **sign cut**: s is equal to 0.
- **gap cut**: s is a value in the largest gap in the sorted list of Fiedler vector components.

Other variations have been proposed.

In this paper, we will analyze the spectral method that uses the splitting value that achieves the best ratio cut. We will show that, for bounded-degree planar graphs and well-shaped meshes, it always finds a good ratio cut. In fact, it is not necessary to use a Fiedler vector; an approximation will suffice. If a vector $\vec{x}$ that is orthorgonal to the all-ones vector has a small Rayleigh quotient with respect to the Laplacian of G, then $\vec{x}$ can be used to find a good ratio cut of G.

Guattery and Miller [GM95] showed that there exist bounded-degree planar graphs on n vertices with constant-size separators for which spectral bisection and spectral sign cuts give separators that cut $n/3$ edges.

We will show that planar graphs have Fiedler cuts of ratio $O(1/\sqrt{n})$. By Lemma 22, our result implies that a bisector of size $O(\sqrt{n})$ can be found by repeatly finding Fiedler cuts. In Section 8, we extend the results of Guattery and Miller to show that this repeated application of Fiedler cuts is necessary, even for some quite natural graphs. We will show that, for any constant δ in the range $0 < \delta \leq 1/2$, there are natural families of well-shaped two-dimensional meshes that have no Fiedler cut of small ratio that is also a δ-separator. We discuss why these graphs exist as well as why they might fail to arise in practice. Our explanation is not entirely satisfactory and the problem remains of finding a better characterization of the graphs that do arise in practice as well as those for which there is a Fiedler cut that produces a δ-separator of size $O(\sqrt{n})$.

3. The eigenvalues of planar graphs

In this section, we will prove that the Fiedler value of every bounded-degree planar graph is $O(1/n)$. Our proof establishes and exploits a connection between the Fiedler value and geometric embeddings of graphs. We obtain the eigenvalue bound by demonstrating that every planar graph has a "nice" embedding in Euclidean space.

A bound of $O(1/\sqrt{n})$ can be placed on the Fiedler value of any planar graph by combining the planar separator theorem of Lipton and Tarjan [LT79] with the fact that $\lambda_2/2 \leq \phi(G)$. Bounds of $O(1/n)$ on the Fiedler values of planar graphs were previously known for graphs such as regular grids [PSL90], quasi-uniform graphs [GK95], and bounded-degree trees. Bounds on the Fiedler values of regular grids and quasi-uniform graphs essentially follow from the fact that the diameters of these graphs are large (see [Chu89]). Bounds on trees can be obtained from the fact that every bounded-degree tree has a δ-separator of size 1 for some constant δ in the range $0 < \delta < 1/2$ that depends only on the degree. However, in order to estimate the Fiedler value of general bounded-degree planar graphs and well-shaped meshes, we need different techniques.

We denote the standard l_2 norm of a vector $\vec{x}$ in Euclidean space by $\|\vec{x}\| = \sqrt{x^T x}$. We relate the quality of an embedding of a graph in Euclidean space with its Fiedler value by the following lemma:

Lemma 1 (embedding lemma). *Let $G = (V, E)$ be a graph. Then λ_2, the Fiedler value of G, is given by*

$$\lambda_2 = \min \frac{\sum_{(i,j)\in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^{n} \|\vec{v}_i\|^2},$$

where the minimum is taken over vectors $\{\vec{v}_1, \ldots, \vec{v}_n\} \subset R^n$ such that

$$\sum_{i=1}^{n} \vec{v}_i = \vec{0},$$

where $\vec{0}$ denotes the all-zeroes vector.

Remark 2. *While we state this lemma for vectors in R^n, it applies equally well for vectors in R^m for any $m \geq 1$.*

Proof: Because the all-ones vector is the eigenvector of $L(G)$ corresponding to the eigenvalue 0, λ_2 can be characterized by

$$\lambda_2 = \min \frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{\sum_{i=1}^{n} x_i^2},$$

where the minimum is taken over real x_i's such that $\sum_{i=1}^{n} x_i = 0$. The minimum is achieved precisely when $(x_1, \ldots, x_n)$ is an eigenvector.

The embedding lemma now follows from a component-wise application of this fact. Write $\vec{v}_i$ as $(v_{i,1}, \ldots, v_{i,n})$. Then, for all $\{\vec{v}_1, \ldots, \vec{v}_n\}$ such that $\sum_{i=1}^{n} \vec{v}_i = \vec{0}$, we have

$$\begin{aligned} \frac{\sum_{(i,j)\in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^{n} \|\vec{v}_i\|^2} &= \frac{\sum_{(i,j)\in E} \sum_{k=1}^{n} (v_{i,k} - v_{j,k})^2}{\sum_{i=1}^{n} \sum_{k=1}^{n} v_{i,k}^2} \\ &= \frac{\sum_{k=1}^{n} \sum_{(i,j)\in E} (v_{i,k} - v_{j,k})^2}{\sum_{k=1}^{n} \sum_{i=1}^{n} v_{i,k}^2}. \end{aligned}$$

But, for each k,

$$\frac{\sum_{(i,j)\in E}(v_{i,k} - v_{j,k})^2}{\sum_{i=1}^{n} v_{i,k}^2} \geq \lambda_2,$$

so

$$\frac{\sum_{k=1}^{n} \sum_{(i,j)\in E}(v_{i,k} - v_{j,k})^2}{\sum_{k=1}^{n} \sum_{i=1}^{n} v_{i,k}^2} \geq \lambda_2$$

(this follows from the fact that $\sum_i x_i / \sum_i y_i \geq \min_i x_i/y_i$, for $x_i, y_i > 0$). □

Our method of finding a good geometric embedding of a planar graph is similar to the way in which Miller, Teng, Thurston, and Vavasis [MTTV96a] directly find good separators of planar graphs.

We first find an embedding of the graph on the plane by using the "kissing disk" embedding of Koebe, Andreev, and Thurston [Koe36, And70a, And70b, Thu88]:

Theorem 3 (Koebe-Andreev-Thurston). *Let G be a planar graph with vertex set $V = \{1, \ldots, n\}$ and edge set E. Then, there exists a set of disks $\{D_1, \ldots, D_n\}$ in the plane with disjoint interiors such that D_i touches D_j if and only if $(i, j) \in E$.*

Such an embedding is called a *kissing disk* embedding of G.

The analogue of a disk on the sphere is a *cap*. A *cap* is given by the intersection of a half-space with the sphere, and its boundary is a circle. We define *kissing caps* analogously with kissing disks. Following [MTTV96a], we use stereographic projection to map the kissing disk embedding of the graph on the plane to a kissing cap embedding on the sphere (See Section 4 for more information on stereographic projection). In Theorem 9, we will show that we can find a sphere preserving map map that sends the *centroid* (also known as the *center of gravity* or *center of mass*) of the centers of the caps to the center of the sphere. Using this theorem, we can bound the eigenvalues of planar graphs:

Theorem 4. *Let G be a planar graph on n nodes of degree at most d. Then, the Fiedler value of G is at most*

$$\frac{8d}{n}.$$

Accordingly, G has a Fiedler cut of ratio $O(1/\sqrt{n})$, and one can iterate Fiedler cuts to find a bisector of size $O(\sqrt{n})$.

Proof: By Theorem 3 and Theorem 9, there is a representation of G by kissing caps on the unit sphere so that the centroid of the centers of the caps is the center of the sphere. Let $\vec{v}_1, \ldots, \vec{v}_n$ be the centers of these caps. Make the center of the sphere the origin, so that $\sum_{i=1}^{n} \vec{v}_i = \vec{0}$.

Let $r_1, \ldots, r_n$ be the radii of the caps. If cap i kisses cap j, then the edge from $\vec{v}_i$ to $\vec{v}_j$ will have length at most $(r_i + r_j)^2$. As this is at most $2(r_i^2 + r_j^2)$, we can divide the contribution of this edge between the two caps. That is, we write

$$\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2 \leq 2d \sum_{i=1}^{n} r_i^2.$$

But, because the caps do not overlap,

$$\sum_{i=1}^{n} \pi r_i^2 \leq 4\pi.$$

Moreover, $\|\vec{v}_i\| = 1$ because the vectors are on the unit sphere.

Applying the embedding lemma, we find that the Fiedler value of G is at most

$$\frac{\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^{n} \|\vec{v}_i\|^2} \leq \frac{8d}{n}.$$

Given the bound on the Fiedler value, the ratio achievable by a Fiedler cut follows immediately from Theorem 21 and the corresponding bisector size follows Lemma 22. $\square$

Remark 5. *One can prove a slightly weaker result by using a result of Miller, Teng, Thurston, and Vavasis [MTTV96a] to find a circle-preserving map that makes the center of the sphere a centerpoint [Ede87] of the images of particular points in the caps. If the center of the sphere is a centerpoint, then the centroid must be far away from at least a constant fraction of the centers of the caps. Thus, the numerator of the Rayleigh quotient will be the same as in Theorem 4, and the denominator will be* $\Omega(n)$*.*

Remark 6. *The embedding lemma can be viewed as a semi-definite relaxation of an integer program for minimum balanced cut. The integer program would be*

$$\begin{aligned} \min \quad & \sum_{(i,j)\in E} (x_i - x_j)^2 \\ s.t. \quad & \sum_{i=1}^{n} x_i = 0, \ \ and \\ & x_i \in \{\pm 1\}. \end{aligned}$$

However, unlike the semi-definite relaxation of MaxCut used by Goemans and Williamson [GW94], we do not know if our relaxation provides a constant factor approximation of the optimum.

4. Sphere-preserving maps

Let B^d be the unit ball in d dimensions: $\{(x_1, \ldots, x_d) | \sum_{i=1}^{n} x_i^2 \leq 1\}$. Let S^d denote the sphere defining the surface of B^d. This section is concerned with *sphere-preserving* maps from S^d to S^d. A *sphere-preserving* map from S^d to S^d is a continuous function that sends every sphere (of lower dimension) contained in S^d to a sphere in S^d and such that every sphere in S^d has a pre-image under the map that is also a sphere. Familiar sphere-preserving maps include rotations and the map that sends each point to its antipode.

We will make use of a slightly larger family of sphere-preserving maps. We obtain this family by first considering sphere-preserving maps between the sphere and the plane. Let H^d be the hyperplane tangent to S^d at $(-1, 0, \ldots, 0)$. One can map H^d to S^d by *stereographic projection*: $\Pi : H^d \to S^d$ by

$$\Pi(z) = \text{the intersection of } S^d \text{ with the line connecting } z \text{ to } (1, 0, \ldots, 0).$$

Similarly, one defines a map $\Pi^{-1} : S^d \to H^d$ that sends a point $z \in S^d$ to the intersection of H^d with the line through z and $(1, 0, \ldots, 0)$. Note that Π^{-1} is not well-defined at $(1, 0, \ldots, 0)$. To fix this, we add the point ∞ to the hyperplane H^d, and define $\Pi^{-1}(1, 0, \ldots, 0) = \infty$ as well as $\Pi(\infty) = (1, 0, \ldots, 0)$.

For any point $\alpha \in S^d$, we define Π_α to be the stereographic projection from the plane perpendicular to S^d at α, and we let Π_α^{-1} be its inverse (so, $\Pi(\infty) = -\alpha$). One can show that the maps Π_α and Π_α^{-1} are sphere-preserving maps (see [HCV52] or [MTTV96a] for a proof).

Sphere-preserving maps in the plane include rigid motions of the plane as well as dilations (and other mobius transformations). We will obtain sphere-preserving maps in the sphere by applying a projection onto a plane, then applying a dilation of the plane, and then mapping back by stereographic projection. Thus, for $\alpha \in S^d$ and $a \geq 0$, we define D^a_α to be the map that dilates the hyperplane perpendicular to S^d at α by a factor of a (note that $D^a_\alpha(\infty) = \infty$). For example,

$$D^a_{(-1,0,\dots,0)} : (-1, x_2, \dots, x_d) \mapsto (-1, ax_2, \dots, ax_d).$$

As the composition of sphere-preserving maps is again a sphere-preserving map, we can now define the sphere-preserving maps that we will use. For any α such that $\|\alpha\| < 1$, define $f_\alpha(z)$ by

$$f_\alpha(z) = \Pi_{\alpha/\|\alpha\|}(D^{1-\|\alpha\|}_{\alpha/\|\alpha\|}(\Pi^{-1}_{\alpha/\|\alpha\|}(z))).$$

It is routine to verify that f_α is continuous. We wish to extend the definition of f_α to α on S^2, even though the resulting maps will not be continuous. For $\|\alpha\| = 1$, we define

$$f_\alpha(z) = \begin{cases} -\alpha & \text{if } z = -\alpha, \text{ and} \\ \alpha & \text{otherwise.} \end{cases}$$

We will now examine the effect of the maps f_α on arrangements of spherical caps on S^d. Recall that a spherical cap on S^d is a connected region of S^d whose boundary is a $(d-1)$-dimensional sphere. Thus, the image of a cap under a map f_α is determined by the image of its boundary along with a point in its interior. For a cap C on S^d, let $p(C)$ denote the point on S^d that is the center of C (i.e., the point inside C that is equidistant from its boundary). We want to show that, for any arrangement of caps $\{C_1, \dots, C_n\}$ on S^d, there is an $\alpha \in S^d$ so that the centroid of $\{p(f_\alpha(C_1)), \dots, p(f_\alpha(C_n))\}$ is the origin. But first, we must exclude some degenerate cases:

Definition 7. An arrangement of caps $\{C_1, \dots, C_n\}$ in S^d is *well-behaved* if there is no point that belongs to at least half of the caps.

Remark 8. *All of the arrangements of caps obtained from graphs contained in the other sections of this paper are well-behaved. Otherwise, the induced graphs would have cliques on half of their vertices and no small separators.*

Theorem 9. *For any well-behaved arrangement of caps $\{C_1, \dots, C_n\}$ in S^d, there is an α so that $\|\alpha\| < 1$ and*

$$\frac{\sum_{i=1}^n p(f_\alpha(C_i))}{n} = \vec{0}.$$

Proof: Consider the map from α to the centroid of $\{p(f_\alpha(C_1)), \ldots, p(f_\alpha(C_n))\}$. We want to show that $\vec{0}$ has a preimage under this map. This would be easier if the map were continuous, but it is not continuous for $\|\alpha\| = 1$: as $-\alpha$ crosses the boundary of C_i, $p(f_\alpha(C_i))$ jumps from one side of the sphere to the other.

To fix this problem, we construct a slightly modified map that is continuous. Because the set of caps is well-behaved, we can choose an $\epsilon > 0$ so that, for all α such that $\|\alpha\| \geq 1 - \epsilon$, most of the caps $\{f_\alpha(C_1), \ldots, f_\alpha(C_n)\}$ are entirely contained within the ball of radius $1/2n$ around $\alpha/\|\alpha\|$. In particular, this implies that f_α does not map the centroid of the centers of the caps to the origin. For $\alpha \in B^d$, we now define the map

$$\phi(\alpha) = \frac{\sum_{i=1}^{n} w(C_i, \alpha) f_\alpha(p(C_i))}{n},$$

where the weight function w is given by

$$w(C, \alpha) = \begin{cases} (2 - d(\alpha, C))/\epsilon & \text{if } d(\alpha, C) \geq 2 - \epsilon, \text{ and} \\ 1 & \text{otherwise,} \end{cases}$$

where by $d(\alpha, C)$, we mean the greatest distance from α to a point in the cap C (for example, if $-\alpha \in C$, then $d(\alpha, C) = 1 + \|\alpha\|$). We have chosen w to be a continuous function of α that goes to zero as $-\alpha$ approaches the boundary of a cap; so, $\phi(\alpha)$ is also a continuous function.

From the fact that $\{C_1, \ldots, C_n\}$ is well-behaved, it is easy to verify that, for $\alpha \in S^d$, $\phi(\alpha)$ lies on the line connecting $\vec{0}$ to α and is closer to α than it is to $-\alpha$. By combining this fact with some elementary algebraic topology, we can use Lemma 10 to show that there is an α such that $\phi(\alpha) = \vec{0}$.

By our choice of ϵ, $\|\alpha\| < 1 - \epsilon$, so all of the terms $w(\alpha, C_i)$ are 1, which implies that f_α is the map that we were looking for. $\square$

Lemma 10. *Let $\phi : B^d \to B^d$ be a continuous function so that, for $\alpha \in S^d$, $\phi(\alpha)$ lies on the line connecting α with $\vec{0}$ and is closer to α than it is to $-\alpha$. Then, there exists an $\alpha \in B^d$ such that $\phi(\alpha) = \vec{0}$.*

Proof: Assume, by way of contradiction, that there is no point $\alpha \in B^d$ such that $\phi(\alpha) = \vec{0}$. Now, consider the map $b(\phi(\alpha))$, where $b : B^d - \left\{\vec{0}\right\} \to S^d$ by

$$b(z) = z/\|z\|.$$

Since b is a continuous map, $b \circ \phi$ is a continuous map of B^d onto S^d that is the identity on S^d. Then $z \mapsto -b(\phi(z))$ is a map from B^d onto S^d that has no fixed point. This contradicts Brouwer's Fixed Point Theorem, which says that every continuous map from B^d into B^d has a fixed point. $\square$

We have shown that, for most collections of balls in H^d, there is a sphere preserving map from H^d to S^d such that the centroid of the centers of the caps is the origin. We now show that one can find such a map by performing a rigid motion of H^d followed by a dilation of H^d followed by stereographic projection.

Definition 11. An arrangement of balls $\{D_1, \ldots, D_n\}$ in H^d is *well-behaved* if there is no point that belongs to at least half of the balls.

Theorem 12. *Let $\{D_1, \ldots, D_n\}$ be a well-behaved collection of balls. Then, there is a point $x \in H^d$ and an $a > 0$ such that the sphere preserving map*

$$g_{x,a} : z \mapsto \Pi(a(z - x))$$

sends the balls to a collection of caps, the centroid of whose centers is the origin.

Proof: [**sketch**] For an $\alpha \in S^d$, consider the map $g_{\Pi^{-1}(\alpha),(1-\|\alpha\|)}$ followed by a rotation of the sphere that sends $(-1, 0, \ldots, 0)$ to α. As we did in the proof of Theorem 9, we can construct a continuous map from α to a weighted centroid of the centers of the caps, which for $\alpha \in S^d$ sends α to a point on the line segment between α and $\vec{0}$. We can then apply Lemma 10 to prove that there is some map α such that the map $g_{\Pi^{-1}(\alpha),(1-\|\alpha\|)}$ sends the centroid of the centers of the caps to the origin. $\Box$

5. The Spectra of k-Nearest Neighbor Graphs

We extend our spectral planar separator theorem to graphs embedded in three and more dimensions. We show that Fiedler cuts of small ratio can be found in α-overlap graphs of k-ply neighborhood systems. One corollary of this extension is that the spectral method finds small ratio cuts for k-nearest neighbor graphs and well-shaped finite element meshes in any fixed dimension. In this section, we analyze intersection graphs and nearest neighbor graphs. Results on overlap graphs and well-shaped meshes will be given in the next section.

In this section and the next, we will use the following notation: We use capital letters to denote balls in R^d. If A is a ball in R^d, then we will use A' to denote its image on the sphere S^{d+1} under stereographic projection. If α is a positive real and A is a ball of radius r, then $\alpha \cdot A$ is the ball with the same center as A and radius αr. Similarly, if A' is a spherical cap of spherical radius r, then $\alpha \cdot A'$ is the spherical cap with the same center as A' and radius αr. Let V_d be the volume of a unit d-dimensional ball. Let A_d be the surface volume of a unit d-dimensional ball.

5.1. Intersection Graphs

The graphs that we consider are defined by neighborhood systems. A *neighborhood system* is a set of closed balls in Euclidean space. A *k-ply* neighborhood system is one in which no point is contained in the interior of more than k of the balls. Given a neighborhood system, $\Gamma = \{B_1, \ldots, B_n\}$, we define the *intersection graph* of Γ to be the undirected graph with vertex set $V = \{B_1, \ldots, B_n\}$ and edge set

$$E = \{(B_i, B_j) : B_i \cap B_j \neq \emptyset\}.$$

For example, the Koebe-Andreev-Thurston embedding theorem says that every planar graph is isomorphic to the intersection graph of some 1-ply neighborhood system in two dimensions.

Let $P = \{p_1, \ldots, p_n\}$ be a point set in $\mathbb{R}^d$. For each $p_i \in P$, let $N_k(p_i)$ be the set of k points closest to p_i in P (if there are ties, break them arbitrarily). A *k-nearest neighbor graph* of P is a graph with vertex set $\{p_1, \ldots, p_n\}$ and edge set

$$E = \{(p_i, p_j) : p_i \in N_k(p_j) \text{ or } p_j \in N_k(p_i)\}.$$

Miller *et. al.* [MTTV96b] show that every k-nearest neighbor graph in $\mathbb{R}^d$ is a subgraph of an intersection graph of a $\tau_d k$-ply neighborhood system, where τ_d is the *kissing number* in d dimensions—the maximum number of nonoverlapping unit balls in $\mathbb{R}^d$ that can be arranged so that they all touch a central unit ball [CS88]. Moreover, the maximum degree of a k-nearest neighbor graph is bounded by $\tau_d k$.

5.2. A Spectral Bound

Theorem 13. *Let G be a subgraph of an intersection graph of a k-ply neighborhood system in $\mathbb{R}^d$ such that the maximum degree of G is Δ. Then, the Fiedler value of $L(G)$ is bounded by $c_d \Delta (k/n)^{2/d}$, where $c_d = 2(A_{d+1}/V_d)^{2/d}$.*

Proof: Let $\Gamma = \{B_1, ..., B_n\}$ be the k-ply neighborhood system of which G is the intersection graph. By Theorem 9, there exits a sphere-preserving map $\phi : \mathbb{R}^d \to S^d$ such that the centroid of the centers of the images of the B_i's is the center of the sphere.

Let $\phi(\Gamma) = \{B_i', ..., B_n'\}$ be the images of the balls in Γ under ϕ. Then, the balls in $\phi(\Gamma)$ also form a k-ply system. Let r_i be the radius of B_i'. Because $V_d r_i^d \leq \text{volume}(B_i')$,

$$\sum_{i=1}^{n} V_d r_i^d \leq \sum_{i=1}^{n} \text{volume}(B_i') \leq k A_{d+1}. \tag{1}$$

By Lemma 1,

$$\begin{aligned} \lambda_2(L(G)) &\leq \frac{\sum_{i=1}^n 2\Delta r_i^2}{n} \\ &\leq (2\Delta) \frac{(k A_{d+1}/V_d)^{2/d} n^{1-2/d}}{n} \\ &\leq (2\Delta) \left(\frac{A_{d+1}}{V_d}\right)^{2/d} \left(\frac{k}{n}\right)^{2/d}. \end{aligned}$$

Note that the second inequality follows from (1). □

The next two corollaries follow from Theorem 13, Theorem 21, and Lemma 22.

Corollary 14. *The Fiedler value of a k-nearest neighbor graph of n points in $\mathbb{R}^d$ is bounded by $O(k^{1+2/d}/n^{2/d})$. Therefore, G has a Fiedler cut of ratio $O\left(k^{1+1/d}/n^{1/d}\right)$, and one can repeatedly take Fiedler cuts to find a bisector of size $O\left(k^{1+1/d} n^{1-1/d}\right)$.*

Corollary 15. *Let G be a subgraph of an intersection graph of a k-ply neighborhood system in R^d whose maximum degree is Δ. Then, G has a Fiedler cut of ratio $O\left(\Delta^{1+1/d}/n^{1/d}\right)$, and one can iterate Fiedler cuts to obtain a bisector of size $O\left(\Delta^{1+1/d}n^{1-1/d}\right)$.*

6. The Spectra of Well-Shaped Meshes

One of the main applications of the spectral method is the partitioning of meshes for parallel numerical simulations. Many experments demonstrate the effectiveness of this method [BS92, HL92, HL93, PSL90, Sim91, Wil90]. In this section, we explain why the spectral method finds such good partitions of well-shaped meshes.

6.1. Well-Shaped Meshes

Most numerical methods work by approximating continuous problems with discrete problems on finite structures whose solutions can be efficiently computed. The finite structure used is often called a *mesh*. Many such methods have been developed and applied to important problems in mechanics and physics.

Most of these numerical methods can be classified as *equation based methods* (e.g., the finite element, finite difference, and finite volume methods) or *particle methods* (e.g., the N-body simulation method). However different the particular methods may be, a basic principle is common to all—accuracy of approximation is ensured by using meshes that satisfy certain numerical and geometric constraints. Meshes that satisfy these constraints are said to be *well-shaped*.

To motivate our spectral analysis of well-shaped meshes, we review the conditions required of finite element and finite difference meshes. More detailed discussions can be found in several books and papers (for example, see [SF73, Joh92, BEG94, BE92, Fri72]). Background material on the particle method can be found in [BH86, GR87, HE81, Zha87].

The *finite element method* approximates a continuous problem by subdividing the domain (a subset of R^d) of the problem into a *mesh* of polyhedral *elements* and then approximates the continuous function by piecewise polynomial functions on the elements. A common choice for an element is a d-dimensional *simplex*. Accordingly, a d-dimensional finite element mesh is a d-dimensional *simplicial complex*, a collection of d-dimensional simplices that meet only at shared faces [BEG94, BE92, MT90].

The computation graph associated with each simplicial complex is often its 1-skeleton or the 1-skeleton of its geometric dual (as used in the finite volume method). In the finite element method, a linear system is defined over a mesh, with variables representing physical quantities at the nodes. The nonzero structure of the coefficient matrix of such a linear system is exactly the adjacency structure of the 1-skeleton of the simplicial complex.

To ensure accuracy, in addition to the conditions that a mesh must conform to the boundaries of the region and be fine enough, each individual element of the mesh must be *well-shaped*. A common shape criterion for the finite element method is that the angles of each element are not too small, or the *aspect ratio* of each element is bounded [BA76, BEG94, Fri72]. Other numerical formulations require slightly different conditions. For example, the controlled volume formulation [Nic92, MTTW95] using a Voronoi diagram requires that the *radius aspect ratio* (the ratio of the circumscribed radius to the shortest edge length of an element in the dual Delaunay diagram) is bounded.

The finite difference method also uses a discrete structure, *a finite difference mesh*, to approximate a continuous problem. Finite difference meshes are often produced by inserting a uniform grid from R^2 or R^3 into the domain via a boundary-matching conformal mapping. Notice that, unlike a finite element mesh, a finite difference mesh need not be a collection of simplices or elements, so we can not analyze it as we do a triangulation. In general, the derivative of the conformal transformation must vary gradually with respect to the mesh size in order to produce good results (See, for example [TWM85]). This means that the mesh will probably satisfy a *density condition* [BB87, MV91].

Let G be an undirected graph and let π be an embedding of its nodes in R^d. We say π is an embedding of *density* α if the following inequality holds for all vertices v in G: Let u be the node closest to v. Let w be the node farthest from v that is connected to v by an edge. Then

$$\frac{||\pi(w) - \pi(v)||}{||\pi(u) - \pi(v)||} \leq \alpha.$$

In general, G is an *α-density graph* in R^d if there exists an embedding of G in R^d with density α.

6.2. Modeling Well-Shaped Meshes

We will use the overlap graph to model well-shaped meshes (Miller *et al* [MTTV96a]). An overlap graph is based on a k-ply neighborhood system. The neighborhood system and a parameter, $\alpha \geq 1$, define an overlap graph: Let $\alpha \geq 1$, and let $\Gamma = \{B_1, \ldots, B_n\}$ be a k-ply neighborhood system in R^d. The *α-overlap graph* of Γ is the graph with vertex set $\{B_1, \ldots, B_n\}$ and edge set

$$\{(B_i, B_j) : (B_i \cap (\alpha \cdot B_j) \neq \emptyset) \text{ and } ((\alpha \cdot B_i) \cap B_j \neq \emptyset)\},$$

where by $\alpha \cdot B$, we mean the ball whose center is the same as the center of B and whose radius is larger by a multiplicative factor of α.

Overlap graphs are good models of well-shaped meshes because each well-shaped mesh in two, three, or higher dimensions is a *bounded-degree subgraph* of some overlap graph (for suitable choices of the parameters α and k). For example,

- Let M be a finite element mesh embedded in R^d in which every element has aspect ratio bounded by a. Then, there is a constant α depending only on d and a so that the 1-skeleton of M is a subgraph of an α-overlap graph of a 1-ply neighborhood system. Moreover its maximum degree is bounded by a constant that also depends only on d and a ([MTTV96a]).

- Let M be a Voronoi diagram (from a finite volume method) in R^d in which the radius aspect ratio of its dual Delaunay diagram is bounded by a. Then there is a constant α depending only on d and a so that the dual Delaunay diagram is an α-density graph ([MTTW95]).

- If G is an α-density graph in R^d, then the maximum degree of G is bounded by a constant depending only on α and d; and, G is a subgraph of an α-overlap graph of a 1-ply neighborhood system ([MV91, MTTV96a]).
- The computation/communication graph used in hierarchical N-body simulation methods (such as the Barnes-Hut's treecode method [BH86] and the fast-multipole method [GR87]) is a subgraph of an α-overlap graph of an $O(\log n)$-ply neighborhood system ([Ten96]).

6.3. Spherical Embeddings of Overlap Graphs

In this section, we show that an α-overlap graph is a subgraph of the intersection graph obtained by projecting its neighborhoods onto the sphere and then dilating each by an $O(\alpha)$ factor. By choosing the proper projection, we are able to use this fact to bound the eigenvalues of these graphs.

Theorem 16. *Let $\alpha \geq 1$. Let A and B be balls in R^d such that*

$$(A \cap (\alpha \cdot B) \neq \emptyset) \text{ and } ((\alpha \cdot A) \cap B \neq \emptyset).$$

Then, $(\pi\alpha + \alpha + \pi) \cdot A'$ touches $(\pi\alpha + \alpha + \pi) \cdot B'$.

Our proof uses two lemmas that handle orthogonal special cases.

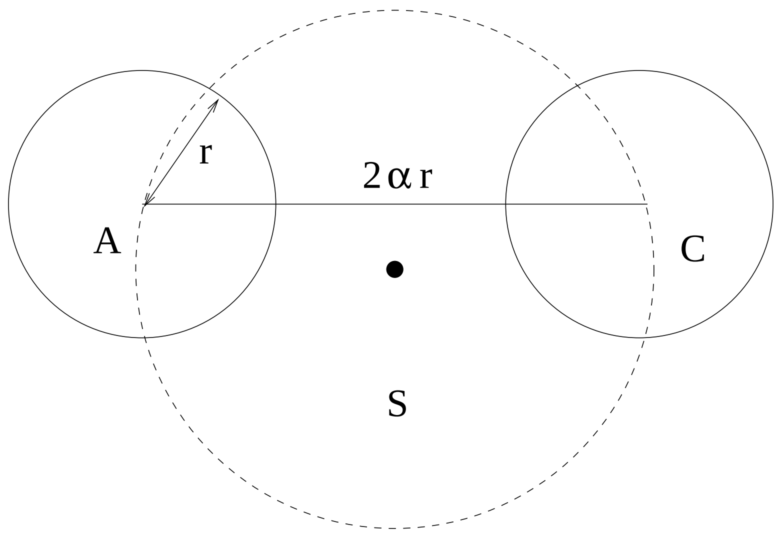

Figure 1:

Lemma 17. *Let A and C be balls in R^d equidistant from the origin and having the same radius. Let A' and C' be their images under stereographic projection onto S^{d+1}. If $\alpha \cdot A$ touches $\alpha \cdot C$, then $(\alpha\pi/2) \cdot A'$ touches $(\alpha\pi/2) \cdot C'$.*

Proof: Let r be the radius of A and C. Because $\alpha \cdot A$ touches $\alpha \cdot C$, the centers of A and C are at distance at most $2\alpha r$ from each other.

Let S be the sphere centered at the origin that passes through the centers of A and C. The geodesic arc between the centers of A and C (on S) has length at most $2\alpha r\pi/2$. The portion of this arc that lies in the interior of A has length at least r (See Figure 1 for a two dimensional example). Since stereographic projection preserves the relations between the intersections of A and C with S, $(\alpha\pi/2) \cdot A'$ will touch $(\alpha\pi/2) \cdot C'$. □

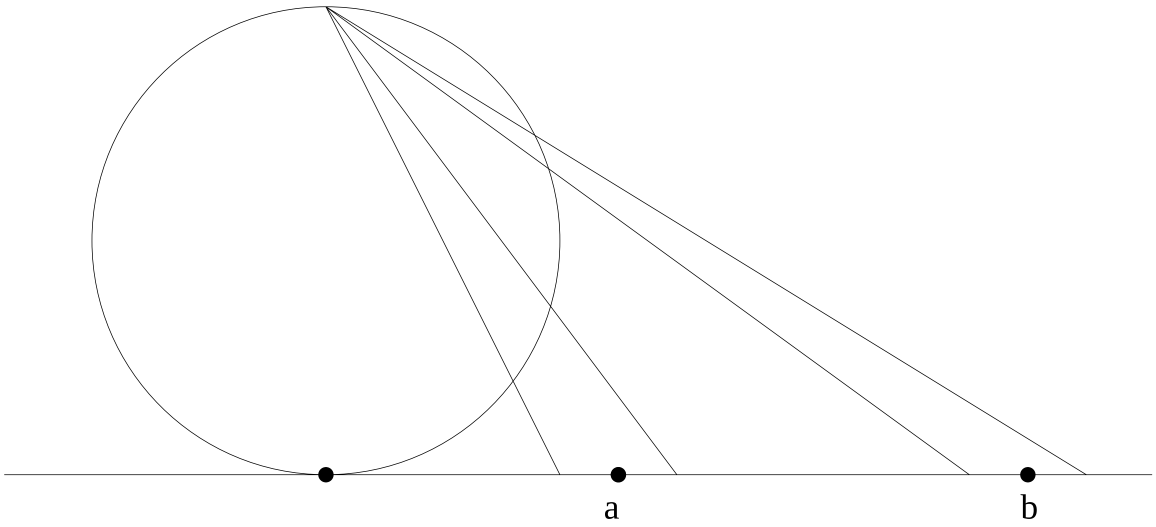

Figure 2: Restriction to the plane through the top of the sphere, the origin, and the centers of A and B.

Lemma 18. *Let A and B be balls in R^d so that the center of A, the center of B, and the origin are colinear and the origin does not lie on the line segment between the center of A and the center of B. If A is closer to the origin than B and $\alpha \cdot A$ touches B, then $\alpha \cdot A'$ touches B'.*

Proof: We will restrict our attention to the plane through the top of the sphere, the origin, and the centers of A and B (see Figure 2). Let a denote the interval that is the intersection of A with the plane. Observe that an interval of the same size as a but located further to the right on the line will have a smaller projection on the circle. The lemma follows. □

Proof: [**of Theorem 16**] Let A and B be any two balls in R^d and let A' and B' be their images under stereographic projection on S^{d+1}. Assume that $\alpha \cdot A$ touches B and $\alpha \cdot B$ touches A. We will show that $(\pi\alpha + \alpha + \pi) \cdot A'$ touches $(\pi\alpha + \alpha + \pi) \cdot B'$.

Assume, without loss of generality, that A is at least as large as B. Let C be the disk of the same distance to the origin as A and congruent to A that is closest to B. Then, the centers of C and B are colinear with the origin (See Figure 3). Let C' be the image of C. Since C is closer to B than A is, $\alpha \cdot C$ touches B and $\alpha \cdot B$ touches A. By Lemma 18, $\alpha \cdot C'$ touches $\alpha \cdot B'$.

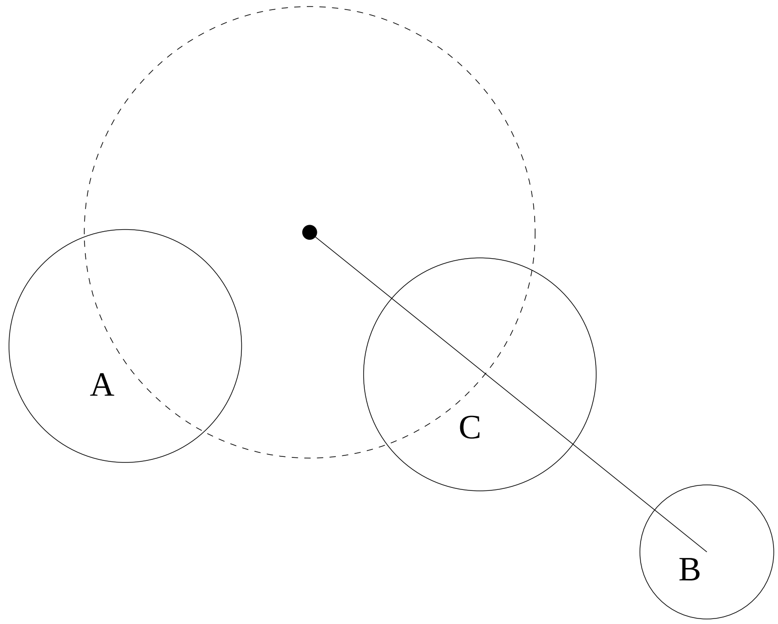

Figure 3: C is the circle congruent to and equidistant from the center to A that is closest to B.

The distance between the centers of A and B is less than $(\alpha+1)$ times the radius of A (because we assume that A is at least as large as B). The same holds for the distance between the center of C and the center of B. Therefore, $(\alpha+1)\cdot A$ touches $(\alpha+1)\cdot C$, so Lemma 17 implies that $\pi(\alpha+1)/2\cdot A'$ touches $\pi(\alpha+1)/2\cdot C'$. Since A' and C' have the same spherical radius, $\alpha\cdot C' \subset (\pi(\alpha+1)+\alpha)A'$. Thus, $(\pi\alpha+\alpha+\pi)\cdot A'$ must touch $(\pi\alpha+\alpha+\pi)\cdot B'$. $\square$

6.4. A Spectral Bound

We now show that the Fiedler value of a bounded degree subgraph of an α-overlap graph is small.

Theorem 19. *If G is a subgraph of an α-overlap graph of a k-ply neighborhood system in R^d and the maximum degree of G is Δ, then the Fiedler value of $L(G)$ is bounded by $\gamma_d\Delta\alpha^2(k/n)^{2/d}$, where $\gamma_d = 2(\pi+1+\pi/\alpha)^2(A_{d+1}/V_d)^{2/d}$. Accordingly, G has a Fiedler cut of ratio $O\left(\Delta\alpha(k/n)^{1/d}\right)$, and one can iterate Fiedler cuts to obtain a bisector of size $O\left(\Delta\alpha k^{1/d}n^{1-1/d}\right)$.*

Proof: Let $\Gamma = \{B_1, ..., B_n\}$ be the k-ply neighborhood system whose intersection graph contains G. By Theorem 12, there is a stereographic projection Π from R^d onto a particular sphere S^{d+1} so that the centroid of the centers of the images of the neighborhoods is the center of the sphere.

Let $\Pi(\Gamma) = \{B_i', ..., B_n'\}$ be the images of the balls in Γ under Π. Let r_i be the radius of B_i'. Because $V_d r^d \leq volume(B_i')$, We know that

$$\sum_{i=1}^{n} V_d r_i^d \leq \sum_{i=1}^{n} volume(B_i') \leq kA_{d+1}.$$

By Theorem 16, G is a subgraph graph of the intersection graph of $\{(\pi\alpha+\alpha+\pi)\cdot B'_i : 1 \leq i \leq n\}$. Thus, by Lemma 1,

$$\begin{aligned}
\lambda_2(L(G)) &\leq \frac{\sum_{i=1}^{n} 2\Delta(\pi\alpha+\alpha+\pi)^2 r_i^2}{n} \\
&\leq (2\Delta)(\pi\alpha+\alpha+\pi)^2 \left(\frac{A_{d+1}}{V_d}\right)^{2/d} \left(\frac{k}{n}\right)^{2/d}.
\end{aligned}$$

Given the bound on the Fiedler value, the ratio achievable by a Fiedler cut follows immediately from Theorem 21 and the corresponding bisector size follows Lemma 22. □

Remark 20. *Recently, Agarwal and Pach [AP95] and, independently, Spielman and Teng [ST96] gave an elementary proof of the sphere separator theorem of Miller* et al *[MTTV96b] on planar graphs and intersection graphs. However, these proofs do not directly extend to overlap graphs. The relation between overlap graphs and intersection graphs established by Theorem 16 enables us to prove the overlap graph separator theorem using the intersection graph separator theorem. The same reduction also extends the deterministic linear time algorithm for finding a good sphere separator from intersection graphs to overlap graphs [EMT95].*

7. Good ratio cuts

Alon [Alo86] and Sinclair and Jerrum [SJ88] proved that graphs with small Fiedler eigenvalue have a good ratio cut (Alon's theorem actually demonstrates the existence of a small vertex separator). A corollary of an extension of their work by Mihail [Mih89] demonstrates that one can obtain a good ratio cut from any vector with small Rayleigh quotient that is perpendicular to the all-ones's vector (although this is not explicitly stated in her work). In this section, we will present a new proof of Mihail's theorem (see also [AM85, Fil91, DS91, Mih89] and [Moh89] for a tighter bound).

Theorem 21 (Mihail). *Let $G = (V, E)$ be a graph on n nodes of maximum degree d, let Q be its Laplacian matrix, and let ϕ be its isoperimetric number. For any vector $\vec{x} \in R^n$ such that $\sum_{i=1}^{n} x_i = 0$,*

$$\frac{\vec{x}^T Q \vec{x}}{\vec{x}^T \vec{x}} \geq \frac{\phi^2}{2d}.$$

Moreover, there is an s so that the cut $(\{i : x_i \leq s\}, \{i : x_i > s\})$ has ratio at most $\phi^2/(2d)$.

Proof: Assume, without loss of generality, that $x_1 \leq x_2 \leq \cdots \leq x_n$, and consider the embedding of G in R by $\vec{x}$. For $i \leq n/2$, at least ϕi edges must cross over x_i. Similarly, for $i \geq n/2$, at least $\phi(n-i)$ edges must cross x_i. We will use this fact to show that the Rayleigh quotient of $\vec{x}$ cannot be too small.

In our proof, we will deal with the $i \leq n/2$ and the $i \geq n/2$ similarly. To simplify matters, assume that n is odd. To achieve symmetry, we will work instead with the vector $\vec{y}$, where $y_i = x_i - x_{(n+1)/2}$, so that $y_{(n+1)/2} = 0$. In a moment, we will show that

$$\frac{\vec{x}^T Q \vec{x}}{\vec{x}^T \vec{x}} \geq \frac{\vec{y}^T Q \vec{y}}{\vec{y}^T \vec{y}},$$

so it suffices to find a lower bound for the Rayleigh quotient of $\vec{y}$ with respect to Q. Recall that

$$\frac{\vec{x}^T Q \vec{x}}{\vec{x}^T \vec{x}} = \frac{\sum_{(i,j)\in E}(x_i - x_j)^2}{\sum_{i=1}^n x_i^2}.$$

Since $(x_i - x_j) = (y_i - y_j)$, the inequality follows from

$$\begin{aligned}
\vec{y}^T \vec{y} = \sum_{i=1}^n (x_i - x_{(n+1)/2})^2 &= \sum_{i=1}^n x_i^2 - 2x_{(n+1)/2} \sum_{i=1}^n x_i + n x_{(n+1)/2}^2 \\
&= \sum_{i=1}^n x_i^2 + n x_{(n+1)/2}^2 \qquad \text{(recall } \textstyle\sum x_i = 0) \\
&\geq \sum_{i=1}^n x_i^2 = \vec{x}^T \vec{x}.
\end{aligned}$$

So that we can treat the $i < (n+1)/2$ and the $i > (n+1)/2$ independently, we would like to eliminate all edges (i,j) where $i < (n+1)/2 < j$. Actually, we will replace each such edge with two edges: one from i to $(n+1)/2$ and one from $(n+1)/2$ to j. Let $\tilde{E}$ denote this new set of edges. Also, let $\tilde{E}_-$ be those edges $(i,j) \in \tilde{E}$ such that $i, j \leq (n+1)/2$, and let $\tilde{E}_+$ be the others. Because $(y_j - y_i)^2 \geq (y_j - y_{(n+1)/2})^2 + (y_{(n+1)/2} - y_i)^2$, we find

$$\begin{aligned}
\frac{\sum_{(i,j)\in E}(y_i - y_j)^2}{\sum_{i=1}^n y_i^2} &\geq \frac{\sum_{(i,j)\in \tilde{E}}(y_i - y_j)^2}{\sum_{i=1}^n y_i^2} \\
&= \frac{\sum_{(i,j)\in \tilde{E}_-}(y_i - y_j)^2 + \sum_{(i,j)\in \tilde{E}_+}(y_i - y_j)^2}{\sum_{i=1}^{(n+1)/2} y_i^2 + \sum_{i=(n+1)/2}^n y_i^2}.
\end{aligned}$$

We now prove a lower bound on the terms involving $i \leq (n+1)/2$. As a similar bound will hold for the other terms, the combination of the two will prove the theorem.

The difficulty in proving a bound on the Rayleigh quotient of $\vec{y}$ is that the numerator is composed of terms like $(y_i - y_j)^2$, while the denominator is a sum of y_i^2's. To overcome this problem, we would like to bound the terms $(y_i - y_j)^2$ by a combination of y_i^2 and y_j^2. Since y_i and y_j are usually separated, we can find a bound that usually works. For any $a > 0$, the inequality

$$(y_i - y_j)^2 \geq \frac{y_i^2 - (1+a)y_j^2}{1 + 1/a},$$

follows from

$$y_i^2 = (y_i - y_j + y_j)^2 \leq (1 + 1/a)(y_i - y_j)^2 + (1+a)y_j^2.$$

We will choose a value for a later in the proof. For now, let d_i^+ be the number of edges (i,j) such that $j > i$, let d_i^- be the number such that $j < i$, and apply the inequality whenever $i < j$ to obtain

$$\frac{\sum_{(i,j)\in\tilde{E}_-}(y_i - y_j)^2}{\sum_{i=1}^{(n+1)/2} y_i^2} \geq \frac{\sum_{(i,j)\in\tilde{E}_-}(d_i^+ - d_i^-(1+a))y_i^2}{(1+1/a)\sum_{i=1}^{n/2} y_i^2}.$$

We will now choose a value of a so that

$$\frac{\sum_{(i,j)\in\tilde{E}_-}(d_i^+ - d_i^-(1+a))y_i^2}{\sum_{i=1}^{n/2} y_i^2} \geq \phi/2.$$

Because the terms y_i^2 are strictly decreasing, it suffices choose a so that

$$\frac{\sum_{i=1}^{k}(d_i^+ - d_i^-(1+a))}{k} \geq \phi/2,$$

for all $1 \leq k \leq n/2$. Since $\sum_{i=1}^{k} d_i^+$ is the number of edges with an endpoint less than y_k and $\sum_{i=1}^{k} d_i^-$ is the number of edges both of whose endpoints are less than y_k, their difference is the number of edges that cross y_k, which is at least ϕk. Moreover, since the maximum degree of a node is d, $\sum_{i=1}^{k} d_i^- \leq (d - \phi k)/2$. Thus, a sufficient constraint on a is that

$$\frac{\phi k - a(d-\phi)k/2}{k} \geq \phi/2 \Rightarrow a \leq \phi/(d-\phi).$$

If we set $a = \phi/(d-\phi)$, then $(1+1/a) = d/\phi$, and we obtain a lower bound on the quotient of

$$\frac{\phi^2}{2d}.$$

□

We now explain how to use good ratio cuts to produce a bisection of a graph.

Lemma 22. *Assume that we are given an algorithm that will find a cut of ratio at most $\phi(k)$ in every k-node subgraph of G, for some monotonically decreasing function ϕ. Then repeated application of this algorithm can be used to find a bisection of G of size at most*

$$\int_{x=1}^{n} \phi(x)dx.$$

Proof: The following algorithm (see [LT79, Gil80]) will find the bisection.

i. Initially, let $D^{(0)} = G$, let A and B be empty sets, and let $i = 0$.

ii. If $D^{(i)}$ is empty, then return A and B; otherwise repeat

  (a) Find a cut of ratio at most $\phi(|D^{(i)}|)$ that divides $D^{(i)}$ into $F^{(i)}$ and $\overline{F^{(i)}}$. We assume that $|F^{(i)}| \leq |\overline{F^{(i)}}|$.

(b) If $|A| \leq |B|$, let $A = A \cup F^{(i)}$; otherwise, let $B = B \cup F^{(i)}$;

(c) Let $D^{(i+1)} = \overline{F^{(i+1)}}$, let $i = i + 1$, and return to step (a).

We assume that the algorithm terminates after t iterations. To show that this algorithm produces a bisection, we need to prove that, for all i in the range $0 \leq i < t$, $\min(|A|, |B|) + |F^{(i)}| \leq n/2$. Because $|F^{(i)}| \leq |\overline{F^{(i)}}|$,

$$\min(|A|, |B|) + |F^{(i)}| \leq (|A| + |B| + |F^{(i)}| + |\overline{F^{(i)}}|)/2 = n/2.$$

We now analyze the total cut size. Because the algorithm finds cuts of ratio at most $\phi(|D^{(i)}|)$ at the ith iteration, the number of edges we cut to separate $F^{(i)}$ is at most

$$\begin{aligned}
\phi(|D^{(i)}|)|F^{(i)}| &= \sum_{j=1}^{|F^{(i)}|} \phi(|D^{(i)}|) \\
&= \sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(|D^{(i)}|) \\
&\leq \sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(j)
\end{aligned}$$

The inequality follows from the fact that ϕ is monotonicly decreasing. The total number of edges cut by this algorithm is at most

$$\begin{aligned}
\sum_{i=0}^{t-1} \phi(|D^{(i)}|)|F^{(i)}| &\leq \sum_{i=0}^{t-1} \left(\sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(j) \right) \\
&= \sum_{j=1}^{n} \phi(j) \\
&\leq \int_1^n \phi(x)dx
\end{aligned}$$

The last inequality follows from the assumption that ϕ is monotonically decreasing. $\square$

Remark 23. *If $\phi(x) = x^{-1/d}$ then*

$$\int_1^n \phi(x)dx = \frac{d}{d-1}(n^{1-1/d} - 1).$$

Lipton and Tarjan [LT79] showed that by repeatedly applying an α-separator of size $\beta\sqrt{n}$, one can obtain a bisection of size $\beta/(1 - \sqrt{1-\alpha})\sqrt{n}$. Gilbert [Gil80] extended this result to graphs with positive vertex weights at the expense of a $1/(1 - \sqrt{2})$ factor in the bisection bound. Djidjev and Gilbert [DG92] further generalized this result to graphs with arbitrary weights. Leighton and Rao [LR88] showed that one can obtain an $O(\alpha)$-approximation to a 1/3-separator from an α-approximation to a ratio cut.

8. When Small Fiedler Cuts Are Unbalanced

Recall that a Fiedler cut of a graph is obtained by taking its Fiedler vector $(x_1, \ldots, x_n)$ and a splitting value s and cutting the edges that cross s. In this section, we present examples of planar graphs for which there is no balanced Fiedler cut of good ratio. In fact, we show that no eigenvector corresponding to a small eigenvalue has a splitting value that induces a good balanced cut. These graphs are generalizations of graphs constructed by Guattery and Miller [GM95]. The properties that we demand of these graphs are easy to achieve, and the reader should be able to generalize our techniques to construct examples in many different contexts. Recently, Guattery and Miller [GM96] have extended our extensions of their work to produce bounded-degree planar graphs in which the ratio achieved by any Fiedler cut is far from the graph's isoperimetric number.

All of our example graphs have cuts of ratio $O(1/\sqrt{n})$ that separate $o(n)$ nodes. This is necessary. In Section 3, we proved that every planar graph has a splitting value that induces a cut of ratio $O(1/\sqrt{n})$. Thus, if there is no unbalanced Fiedler cut of ratio $O(1/\sqrt{n})$, then the Fiedler cut of ratio $O(1/\sqrt{n})$ must be balanced. This might explain why most graphs that arise in practice have balanced Fiedler cuts: a fairly regular planar graph should not have a cut of ratio $O(1/\sqrt{n})$ that separates few vertices.

The graphs that we build in this section will be constructed from:

- P_k, the path graph on k vertices ($V = \{0, \ldots, k-1\}$ and $E = \{(i-1, i) : 1 \leq i < k\}$),
- R_k, the ring graph on k vertices ($V = \{0, \ldots, k-1\}$ and $E = \{(i, i+1 \bmod k) : 1 \leq i \leq k\}$),
- $W_{i,j}$, the Cartesian product of P_i with R_j, and
- B_i, the complete binary tree with 2^i leaves and a total of $2^{i+1} - 1$ nodes.

Proposition 24. *The Fiedler value of P_k (the path graph of k vertices) is equal to*

$$4 \sin^2 \left(\frac{\pi}{2k} \right).$$

The Fiedler vector $\vec{u}$ is given by

$$\vec{u}_i = \cos \left(\frac{(2i-1)\pi}{2k} \right).$$

The s-th eigenvalue of P_k is equal to

$$4 \sin^2 \left(\frac{(s-1)\pi}{2k} \right).$$

Its eigenvector $\vec{u}$ is given by

$$\vec{u}_i = \cos \left(\frac{(2i-1)(s-1)\pi}{2n} \right).$$

Proposition 25. *The Fiedler value of R_k is equal to*

$$4\sin^2\left(\frac{\pi}{k}\right).$$

A Fiedler vector $\vec{u}$ is given by

$$\vec{u}_i = \cos\left(\frac{2i\pi}{k}\right)$$

Proposition 26. *Let G be the Cartesian product of graphs G_1 and G_2. Then the eigenvalues of G are equal to all the possible sums of eigenvalues of G_1 and G_2. Therefore, the Fiedler value of G is the smaller of the Fiedler values of G_1 and G_2.*

We now consider graphs $S_{k,a,b}$ obtained by joining one copy of $W_{a,b}$ with two copies of P_k, which we label P_k and P'_k. Label the vertices of $W_{a,b}$ by $\{w_{i,j}\}$, for $0 \leq i < a$ and $0 \leq j < b$. Similarly, label the vertices of P_k and P'_k by p_i and p'_i for $0 \leq i < k$. The graph $S_{k,a,b}$ contains the edges of $W_{a,b}$, P_k, and P'_k, as well as edges connecting $w_{0,i}$ to p_{k-1} and $w_{a-1,i}$ to p'_{k-1}, for all $0 \leq i < k$ (See Figure 4).

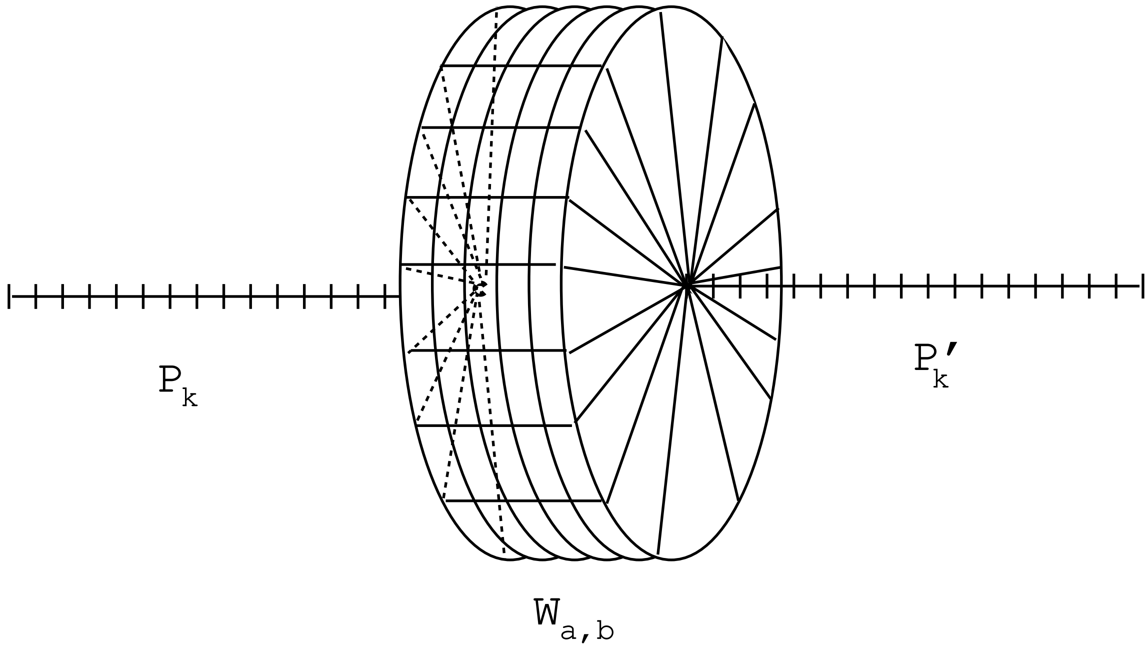

Figure 4: $S_{k,a,b}$: despite the picture, the graph is planar.

We now examine what a Fiedler vector of $S_{k,a,b}$ looks like. We will write a Fiedler vector of $S_{k,a,b}$ as $(\vec{p}, \vec{w}, \vec{p'})$.

Proposition 27. *Let $\vec{u} = (\vec{p}, \vec{w}, \vec{p'})$ be a Fiedler vector of $S_{k,a,b}$. If $k > b > a$, then it cannot be the case that $\vec{p} = \vec{p'} = \vec{0}$.*

Proof: If $\vec{p} = \vec{p'} = \vec{0}$, then the Fiedler value of $S_{k,a,b}$ is at least the Rayleigh quotient of $\vec{w}$ with respect to the Laplacian of $W_{a,b}$: all the terms on the bottom of the Rayleigh quotient of $\vec{u}$ with respect to $L(S_{k,a,b})$ appear on the bottom, and some extra terms appear on the top corresponding to the edges between $\left\{p_{k-1}, p'_{k-1}\right\}$ and $W_{a,b}$. This would imply that the Fiedler value of $S_{k,a,b}$ is at least the Fiedler value of $W_{a,b}$, which in turn larger than the Fiedler value of P_{2k-1}. This is a contradiction because the Fiedler value of P_{2k-1} is an upper bound on the Fiedler value of $S_{k,a,b}$: construct a vector in which the terms corresponding to p_{k-1}, p'_{k-1}, and $W_{a,b}$ are zero, and the remaining nodes are set to the values of the Fiedler vector of P_{2k-1}. $\square$

Proposition 28. *Let $(\vec{p}, \vec{w}, \vec{p'})$ be a Fiedler vector of $S_{k,a,b}$, for $k > b > a$. Then $w_{i,j} = w_{i,k}$ for all $1 \leq j < k \leq b$.*

Proof: Consider the automorphism τ of $S_{k,a,b}$ that maps $w_{i,j}$ to $w_{i,(j+1 \bmod b)}$ and leaves P_k and P'_k fixed. Assume, by way of contradiction, that $\tau(\vec{u}) - \vec{u} \neq \vec{0}$. Then, $\tau(\vec{u}) - \vec{u}$ is a Fiedler vector in which the paths P_k and P'_k get mapped to zero. This would contradict Proposition 27. $\square$

Theorem 29. *Let $k > b > a$. Then, any Fiedler cut of the $(ab + 2k)$-node graph $S_{k,a,b}$ either cuts at least b edges or separates fewer than $2k$ vertices.*

Proof: Immediate from Propositions 27 and 28. $\square$

We now prove a similar statement for a bounded-degree planar graph. We do this by replacing the star graph in $S_{k,a,b}$ with a complete binary tree. For b a power of two, we define the graph $T_{k,a,b}$ to be a graph with two copies of P_k, which we call P_k and P'_k, two copies of B_b, which we call B_b and B_b', and one copy of $W_{a,b}$. These graphs are linked by identifying the leaves of B_b with the nodes $\{w_{0,i}\}_i$ and the leaves of B'_b with the nodes $\{w_{a,i}\}_i$ so that the graph is symmetric, the rightmost leaf of B_b maps to $w_{0,0}$, and the graph remains planar. We attach P_k and P'_k by identifying p_{k-1} with the root of B_b and p'_{k-1} with the root of B'_b (See Figure 5).

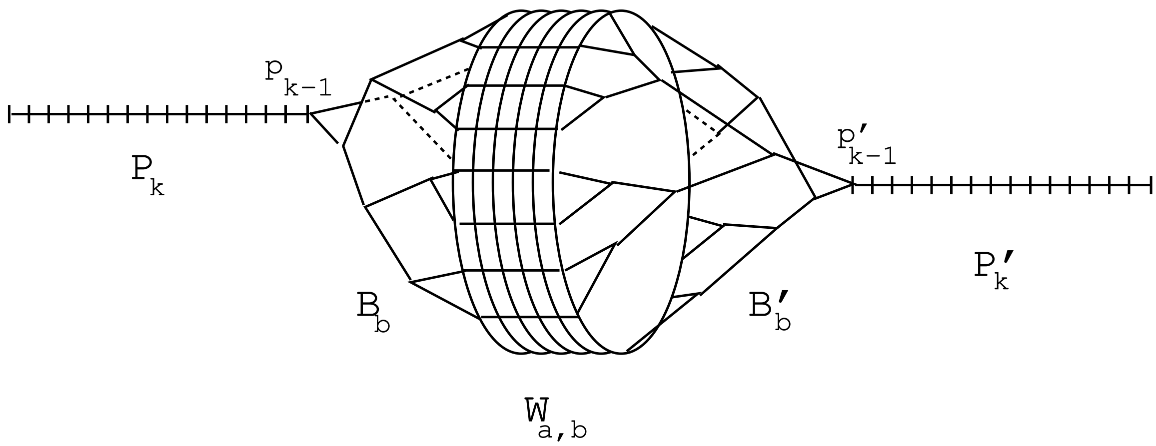

Figure 5: $T_{k,a,b}$

In order to prove a statement similar to Proposition 27, we need to lower bound the Fiedler value of $T_{0,a,b}$. We do this by proving a lower bound on the isoperimetric number of $T_{0,a,b}$ and then applying Theorem 21.

Proposition 30. *For $a \leq b$,*

$$\phi(T_{0,a,b}) \geq 1/(b+1).$$

Proof: Let $(A, \bar{A})$ be a cut of $T_{0,a,b}$ in which $|A| \leq \left|\bar{A}\right|$. Assume that r nodes of A are in the portion of $T_{0,a,b}$ that is $W_{a,b}$, and that t nodes of A are in the internal nodes of the trees B_b and B'_b. Let

$$I = \{i : \exists j \text{ for which } w_{i,j} \in A\}, \text{ and}$$

$$J = \{j : \exists i \text{ for which } w_{i,j} \in A\}.$$

If $|I| < a$, then for each $j \in J$, there is an i such that $(w_{i,j}, w_{i+1,j}) \in E(A, \bar{A})$. Similarly, if $|J| < b$, then for each $i \in I$, there is an i such that $(w_{i,j}, w_{i,j+j}) \in E(A, \bar{A})$. Thus, $\left|E(A, \bar{A})\right|$ is at least the minimum of a, b, and $\max\{|I|, |J|\}$. Since, $r \leq |I| \cdot |J|$, $\left|E(A, \bar{A})\right|$ is at least the minimum of a, b, and $\sqrt{r}$.

The number of edges that leave the t vertices contained in the trees is at least t. Unless $|I| = a$, an edge must be cut for each of these edges. If $|I| = a$, then $E(A, \bar{A})/|A| \geq 1/(1+b)$. Otherwise, $\left|E(A, \bar{A})\right|$ is at least the maximum of t and $\min(a, \sqrt{r})$. Since $|A|$ is $t + r$, we find

$$\frac{E(A, \bar{A})}{|A|} \geq \frac{\max(t, \min(\sqrt{r}, a))}{t + r} \geq 1/(1+b).$$

□

Corollary 31. *The Fiedler value of $T_{0,a,b}$ is at least $\frac{1}{8(b+1)^2}$.*

Proof: Apply Theorem 21, observing that the maximum degree of $T_{0,a,b}$ is 4. □

Proposition 32. *Let $\vec{u} = (\vec{p}, \vec{b}, \vec{w}, \vec{b}', \vec{p}')$ be a Fiedler vector of $T_{k,a,b}$. If $k/2\pi > b > a$, then it cannot be the case that $\vec{p} = \vec{p}' = \vec{0}$.*

Proof: If $\vec{p} = \vec{p}' = \vec{0}$, then the Fiedler value of $T_{k,a,b}$ is at least the Rayleigh quotient of $(\vec{b}, \vec{w}, \vec{b}')$ with respect to the Laplacian of $T_{0,a,b}$: all the terms on the bottom of the Rayleigh quotient of $\vec{u}$ with respect to $L(T_{k,a,b})$ appear on the bottom, and some extra terms appear on the top corresponding to the edges between $\left\{p_{k-1}, p'_{k-1}\right\}$ and B_b and B'_b. This would imply that the Fiedler value of $T_{k,a,b}$ is at least the Fiedler value of $T_{0,a,b}$, which in turn larger than the Fiedler value of P_{2k-1}. This is a contradiction because the Fiedler value of P_{2k-1} is an upper bound on the Fiedler value of $T_{k,a,b}$. We can construct a vector in which the terms corresponding to p_{k-1}, p'_{k-1}, B_b, B'_b and $W_{a,b}$ are zero, and the remaining nodes are set to the values of the Fiedler vector of P_{2k-1}. □

Proposition 33. *Let $\vec{u} = (\vec{p}, \vec{b}, \vec{w}, \vec{b'}, \vec{p'})$ be a Fiedler vector of $T_{k,a,b}$ for $k/2\pi > b > a$. Then $w_{i,j} = w_{i,k}$ for all $1 \le j < k \le b$. Moreover the value assigned to a node in one of the trees only depends on its height in the tree.*

Proof: We first show that the children of the roots of the trees must have the same value. Consider the automorphism of $T_{k,a,b}$ induced by swapping $w_{i,j}$ with $w_{i,b-j-1}$. By an argument similar to that used in the proof of Proposition 28, we see that the values assigned by the Fiedler vector to $w_{i,j}$ and $w_{i,b-j-1}$ must be identical, and the same holds for the tree nodes paired by the mapping.

We can now work our way down the tree. The maps that we use will not necessarily be automorphisms, but they will take advantage of the identification of values that we have established before. Next, we consider the mapping induced by swapping $w_{i,j}$ with $w_{i,b/2-j-1}$, for $0 \le j < b/2$, and $w_{i,b/2+j}$ with $w_{i,b-j-1}$, for $0 \le j < b/2$. This must again produce a Fiedler vector. By applying an argument such as that in the proof of Proposition 28, we obtain more identifications of values.

We conclude the proof by continuing this way down the subtrees. □

Theorem 34. *Let $k/2\pi > b > a$. Then, any Fiedler cut of the $(ab+2(b-1)+2(k-1))$-node graph $T_{k,a,b}$ either cuts at least b edges or separates fewer than $2k$ vertices.*

Proof: Immediate from Propositions 32 and 33. □

For sufficiently large k, no eigenvector of a small eigenvalue of $T_{k,a,b}$ will produce a balanced cut of small ratio. To prove this, we show that the columns of $W_{a,b}$ map to the same point in these eigenvectors as well. We rely on the Courant-Fisher characterization of the eigenvalues of a matrix, which is a generalization of Rayleigh's characterization.

Proposition 35 (Courant-Fisher). *Let M be an $n \times n$ real symmetric matrix. Let $\lambda_k(M)$ denote its k-th smallest eigenvalue. Then*

$$\lambda_k = \min_{U} \max_{\vec{x} \in U} \frac{\vec{x}^T M \vec{x}}{\vec{x}^T \vec{x}},$$

where U ranges over all k-dimensional subspaces of R^n.

Corollary 36. *For $0 < i < k$,*

$$\lambda_{i+1}(T_{k,a,b}) \le \lambda_{2i}(P_{2k-1}).$$

Proof: To bound $\lambda_{i+1}(T_{k,a,b})$, we will construct a set of $i+1$ vectors $\vec{x}_j = (\vec{p}_j, \vec{b}_j, \vec{w}_j, \vec{b'}_j, \vec{p'}_j)$, for $1 \le j \le i+1$ such that

$$\max_{\vec{x} \in \mathrm{span}(\vec{x}_1, \ldots, \vec{x}_{j+1})} \frac{\vec{x}^T L(T_{k,a,b}) \vec{x}}{\vec{x}^T \vec{x}} = \lambda_{2i}(P_{2k-1}).$$

Let $\vec{v}_j$ be the eigenvector of P_{2k-1} corresponding to λ_j. Since $\vec{v}_1$ is the all-ones vector, we let $\vec{x}_1$ be the all-ones vector. We now build the other $\vec{x}_j$'s from the even eigenvectors of P_{2k-1} Since $\vec{v}_{2j}$

maps the middle vertex of P_{2k-1} to zero, we can build an analogous vector $\vec{x}_{j+1}$ by setting all of the middle nodes of $T_{k,a,b}$ to zero and making the paths on the end correspond to the values of $\vec{v}_{2j}$. That is, we set $\vec{b}_{j+1}$, $\vec{w}_{j+1}$, and $\vec{b'}_{j+1}$ to be zero vectors, and we then assign values to $\vec{p}_{j+1}$ and $\vec{p'}_{j+1}$ so that the resulting embedding of $T_{k,a,b}$ resembles the embedding of P_{2k-1} under $\vec{v}_{2j}$, except that there is a large mass at zero. We conclude by applying Proposition 35. $\square$

Corollary 37. *Let $k/4\pi(j-1) > b > a$ and $1 \leq i \leq j$. Then, any cut from the i-th eigenvector of $T_{k,a,b}$ either cuts at least b edges or separates fewer than 2k vertices.*

Proof: Follows by generalizing Proposition 32 and Proposition 33 with Corollary 36. $\square$

Before we conclude this section, we wish to point out some ways in which one can vary the construction of $T_{k,a,b}$ without adversely effecting its resistance to Fiedler cuts. First, it is not necessary to have two sets of paths and two sets of trees leaving the middle section: Using only one path will suffice, although it complicates the proof of Proposition 32. Second, it is not necessary to have dangling paths leaving the graph. One can obtain similar constructions in fully triangulated graphs. A nested chain of triangles (Figure 6) will serve in place of a path.

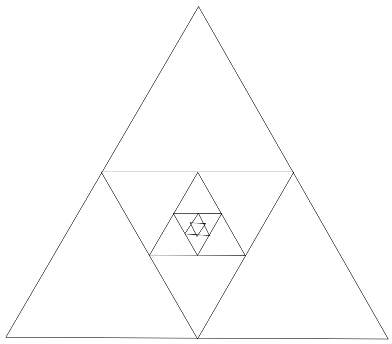

Figure 6: Nested triangles can substitute for path graphs.

9. Acknowledgements

We thank Michael Mandell for advice on algebraic topology, John Gilbert, Nabil Kahale, Gary Miller, and Horst Simon for helpful discussions on the spectra of graphs, and Edmond Chow, Alan Edelman, Steve Guattery, Bruce Hendrickson, Satish Rao, Ed Rothberg, Dafna Talmor, and Steve Vavasis for comments on an early draft of this paper.

References

[AK95] C. J. Alpert and A. B. Kahng. Recent directions in netlist partitioning: A survey. *Integration: The VLSI Journal*, 19:1–81, 1995.

[Alo86] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[AM85] N. Alon and V.D. Milman. λ_1, isoperimetric inequalities for graphs, and superconcentrators. *J. Comb. Theory Series B*, 38:73–88, 1985.

[And70a] E. M. Andreev. On convex polyhedra in lobacevskii space. *Math. USSR Sbornik*, 10(3):413–440, 1970.

[And70b] E. M. Andreev. On convex polyhedra of finite volume in lobacevskii space. *Math. USSR Sbornik*, 12(2):270–259, 1970.

[AP95] P. Agarwal and J. Pach. *Combinatorial Geometry.* Wiley-Interscience, 1995.

[AST90] Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, pages 293–299. ACM, 1990.

[BA76] I. Babuška and A.K. Aziz. On the angle condition in the finite element method. *SIAM J. Numer. Anal.*, 13(2):214–226, 1976.

[Bar82] E. R. Barnes. An algorithm for partitioning of nodes of a graph. *SIAM J. on Algebraic and Discrete Methods*, 4(3):541–550, 1982.

[BB87] M. J. Berger and S. Bokhari. A partitioning strategy for nonuniform problems on multiprocessors. *IEEE Trans. Comp.*, C-36:570–580, 1987.

[BE92] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In F. K. Hwang and D.-Z. Du, editors, *Computing in Euclidean Geometry*, pages 23–90. World Scientific, 1992.

[BEG94] M. Bern, D. Eppstein, and J. R. Gilbert. Provably good mesh generation. *Journal of Computer and System Sciences*, 48:384–409, 1994.

[BH84] E. R. Barnes and A. J. Hoffman. Partitioning, spectra and linear programming. *Progress in Combinatorial Optimization*, pages 13–25, 1984.

[BH86] J. Barnes and P. Hut. A hierarchical $O(n \log n)$ force calculation algorithm. *Nature*, (324):446–449, 1986.

[Big93] Norman Biggs. *Algebraic Graph Theory.* Cambridge University Press, New York, NY, second edition, 1993.

[Bop87] R. Boppana. Eigenvalues and graph bisection: An average-case analysis. In *28th Annual Symposium on Foundations of Computer Science*, pages 280–285, Los Angeles, October 1987. IEEE.

[BS92] S. T. Barnard and H. D. Simon. A fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems. Technical Report RNR–92–033, NASA Ames Research Center, 1992.

[CDS90] Dragos M. Cvetkovic, Michael Doob, and Horst Sachs. *Spectra of Graphs : Theory and Application.* Academic Press, New York, 1990.

[Che70] J. Cheeger. A lower bound for the smallest eigenvalue of the laplacian. In R. C. Gunning, editor, *Problems in Analysis*, pages 195–199. Princeton University Press, 1970.

[Chu89] F.R.K. Chung. Diamaters and eigenvalues. *Journal of the American Mathematical Society*, 2(2):187–196, 1989.

[CR87] T. F. Chan and D. C. Resasco. A framework for the analysis and construction of domain decomposition preconditioners. Technical Report CAM-87-09, UCLA, 1987.

[CS88] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups.* Springer-Verlag, 1988.

[CS93] T. F. Chan and B. Smith. Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes. *Contemporary Mathematics*, pages 1–14, 1993.

[CSZ93] P. K. Chan, M. Schlag, and J. Zien. Spectral k-way ratio cut partitioning and clustering. In *Symp. on Integrated Systems*, 1993.

[DG92] H. N. Djidjev and J. R. Gilbert. Separators in graphs with negative and multiple vertex weights. Technical Report CS-92-07, Xerox PARC, Palo Alto, 1992.

[DH72] W. E. Donath and A. J. Hoffman. Algorithms for partitioning of graphs and computer logic based on eigenvectors of connection matrices. *IBM Technical Disclosure Bulletin*, 15:938–944, 1972.

[DH73] W. E. Donath and A. J. Hoffman. Lower bounds for the partitioning of graphs. *J. Res. Develop.*, 17:420–425, 1973.

[DS91] P. Diaconis and D. Strook. Geometric bounds for eigenvalues of Markov chains. *The Annals of Applied Probability*, 1(1):36–61, 1991.

[Ede87] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, NY, 1987.

[EMT95] D. Eppstein, G. L. Miller, and S.-H. Teng. A deterministic linear time algorithm for geometric separators and its applications. *Fundamenta Informaticae*, 22(4):309–330, April 1995.

[Fie73] M. Fiedler. Algebraic connectibity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.

[Fie75a] M. Fiedler. Eigenvectors of acyclic matices. *Czechoslovak Mathematical Journal*, 25(100):607–618, 1975.

[Fie75b] M. Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its applications to graph theory. *Czechoslovak Mathematical Journal*, 25(100):619–633, 1975.

[Fil91] James Allen Fill. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The Annals of Applied Probability*, 1(1):62–87, 1991.

[Fri72] I. Fried. Condition of finite element matrices generated from nonuniform meshes. *AIAA J.*, 10:219–221, 1972.

[GHT84] J.R. Gilbert, J.P. Hutchinson, and R.E. Tarjan. A separation theorem for graphs of bounded genus. *Journal of Algorithms*, 5:391–407, 1984.

[Gil80] John Russell Gilbert. *Graph Separator Theorems and Sparse Gaussian Elimination.* PhD thesis, Computer Science Department, Stanford University, 1980.

[GK95] J. R. Gilbert and N. Kahale. personal communication, 1995.

[GM95] Stephen Guattery and Gary L. Miller. On the performance of the spectral graph partitioning methods. In *Proceedings of The Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 233–242. ACM-SIAM, 1995.

[GM96] Stephen Guattery and Gary L. Miller. personal communication, 1996.

[GR87] L. Greengard and V. Roklin. A fast algorithm for particles simulations. *J. of Comput. Phy.*, 73:325–348, 1987.

[GW94] Michel X. Goemans and David P. Williamson. .878-approximation algorithms for MAX CUT and MAX 2SAT. In *Proceedings of the -42th Annual ACM Symposium on Theory of Computing*, pages 422–431, 1994.

[HCV52] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination.* Chelsea Publishing Company, New York, 1952.

[HE81] R. W. Hackney and J. W. Eastwood. *Computer Simulation Using Particles.* McGraw Hill, 1981.

[HK92] L. Hagen and A. B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Transactions on Computer-Aided Design*, 11(9):1074–1085, September 1992.

[HL92] Bruce Hendrickson and Robert Leland. An impoved spectral graph partitioning algorithm for mapping parallel computation. Technical Report SAND92-1460, Sandia National Lab., December 1992.

[HL93] Bruce Hendrickson and Robert Leland. The Chaco user's guide, Version 1.0. Technical Report SAND93-2339, Sandia National Lab., 1993.

[Joh92] C. Johnson. *Numerical solution of partial differential equations by the finite element method.* Cambridge University Press, 1992.

[Koe36] Paul Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse*, 88:141–164, 1936.

[LR88] F. T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *29th Annual Symposium on Foundations of Computer Science*, pages 422–431, 1988.

[LT79] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. of Appl. Math.*, 36:177–189, April 1979.

[Mih89] M. Mihail. Conductance and convergence of Markov chains – a combinatorial treatment of expanders. In *30th Annual Symposium on Foundations of Computer Science*, pages 526–531. IEEE, 1989.

[Moh88] B. Mohar. The Laplacian spectrum of graphs. In *Sixth International Conference on the Theory and Applications of Graphs*, pages 871–898, 1988.

[Moh89] B. Mohar. Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B*, 47:274–291, 1989.

[MT90] Gary L. Miller and William Thurston. Separators in two and three dimensions. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, pages 300–309, Baltimore, May 1990. ACM.

[MTTV96a] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Finite element meshes and geometric separators. *SIAM J. Scientific Computing*, page to appear, 1996.

[MTTV96b] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Separators for sphere-packings and nearest neighborhood graphs. submitted to *J. ACM*, 1996.

[MTTW95] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, and Noel Walkington. A Delaunay based numerical method for three dimensions: generation, formulation, and partition. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, pages 683–692, Las Vegas, May 1995. ACM.

[MTV91] G. L. Miller, S.-H. Teng, and S. A. Vavasis. An unified geometric approach to graph separators. In *32nd Annual Symposium on Foundations of Computer Science*, pages 538–547. IEEE, 1991.

[MV91] Gary L. Miller and Stephen A. Vavasis. Density graphs and separators. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 331–336, 1991.

[Nic92] R. A. Nicolaides. Direct discretization of planar div-curl problems. *SIAM J. Numer. Anal.*, 29(1), 1992.

[PRS94] S. Plotkin, S. Rao, and W. D. Smith. Shallow excluded minors and improved graph decomposition. In *5th Symp. Discrete Algorithms*, pages 462–470. SIAM, 1994.

[PSL90] A. Pothen, H. D. Simon, and K.-P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.*, 11:430–452, 1990.

[PSS82] B. Parlett, H. Simon, and L. Stringer. Estimating the largest eigenvalues with the lanczos algorithm. *Mathematics of Computation*, 38:153–165, 1982.

[PSW92] A. Pothen, H. D. Simon, and L. Wang. Spectral nested dissection. Technical Report CS–92–01, Pennsylvania State University Department of Computer Science, 1992.

[SF73] G. Strang and G. J. Fix. *An Analysis of the Finite Element Method.* Prentice-Hall, Englewood Cliffs, New Jersey, 1973.

[Sim91] H. D. Simon. Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering*, 2(2/3):135–148, 1991.

[SJ88] A. J. Sinclair and M. R. Jerrum. Conductance and the rapid mixing property for markov chains: the approximation of the permanent resolved. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 235–244. ACM, 1988.

[SJ89] Alistair Sinclair and Mark Jerrum. Approximative counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, 82(1):93–133, July 1989.

[ST96] D. A. Spielman and S.-H. Teng. Disk packings and planar separators. In *Proceedings of 12th ACM Symposium on Computational Geometry*, page to appear. ACM, May 1996.

[Str88] G. Strang. *Linear Algebra and Its Applications.* Harcourt Brace Jovanovich, Publishers, San Diego, CA, 1988.

[Ten91] S.-H. Teng. *Points, Spheres, and Separators: a unified geometric approach to graph partitioning.* PhD thesis, Carnegie-Mellon University, School of Computer Science, 1991. CMU-CS-91-184.

[Ten96] Shang-Hua Teng. Provably good partitioning and load balancing algorithms for parallel adaptive n-body simulation. *SIAM J. Scientific Computing (to appear)*, 1996.

[Thu88] W. P. Thurston. *The geometry and topology of 3-manifolds.* Princeton University Notes, 1988.

[TWM85] J. F. Thompson, Z. U. A. Warsi, and C. W. Mastin. *Numerical Grid Generation: Foundations and Applications.* New York, North Holland, 1985.

[Ung51] P. Ungar. A theorem on planar graphs. *Journal London Math Soc.*, 26:256–262, 1951.

[Wil90] R. D. Williams. Performance of dynamic load balancing algorithms for unstructured mesh calculations. Technical report, California Institute of Technology, 1990.

[Zha87] F. Zhao. An O(n) algorithm for three-dimensional n-body simulation. Technical Report TR AI Memo 995, MIT, AI Lab., October 1987.